

Basic properties of the infinite critical-FK random map

Linxiao Chen*

Abstract

In this paper we investigate the critical Fortuin-Kasteleyn (cFK) random map model. For each $q \in [0, \infty]$ and integer $n \geq 1$, this model chooses a planar map of n edges with a probability proportional to the partition function of critical q -Potts model on that map. Sheffield introduced the hamburger-cheeseburger bijection which maps the cFK random maps to a family of random words, and remarked that one can construct infinite cFK random maps using this bijection. We make this idea precise by a detailed proof of the local convergence. When $q = 1$, this provides an alternative construction of the UIPQ. In addition, we show that the limit is almost surely one-ended and recurrent for the simple random walk for any q , and mutually singular in distribution for different values of q .

Keywords. Fortuin-Kasteleyn percolation, random planar maps, hamburger-cheeseburger bijection, local limits, recurrent graph, ergodicity of random graphs.

Mathematics Subject Classification (2010). 60D05, 60K35, 05C81, 60F20.

1 Introduction

Planar maps. Random planar maps has been the focus of intensive research in recent years. We refer to [1] for the physics background and motivations, and to [27] for a survey of recent results in the field.

A *finite planar map* is a proper embedding of a finite connected graph into the two-dimensional sphere, viewed up to orientation-preserving homeomorphisms. *Self-loops* and *multiple edges* are allowed in the graph. In this paper we will not deal with non-planar maps, and thus we drop the adjective “planar” sometimes. The *faces* of a map are the connected components of the complement of the embedding in the sphere and the *degree* of a face is the number of edges incident to it. A map is a *triangulation* (resp. a *quadrangulation*) if all of its faces are of degree three (resp. four). The *dual map* \mathbf{M}^\dagger of a planar map \mathbf{M} has one vertex associated to each face of \mathbf{M} and there is an edge between two vertices if and only if their corresponding faces in \mathbf{M} are adjacent.

A *corner* in a planar map is the angular section delimited by two consecutive half-edges around a vertex. It can be identified with an oriented edge using the orientation of the sphere. A *rooted* map is a map with a distinguished oriented edge or, equivalently, a corner. We call *root edge* the distinguished oriented edge, and *root vertex* (resp. *root face*) the vertex (resp. face) incident to the distinguished corner. Rooting a map on a corner (instead of the more traditional choice of rooting on an oriented edge) allows a canonical choice of the root for the dual map: the dual root is obtained by exchanging the root face and the root vertex. A *subgraph* of a planar map is a graph consisting of a subset of its edges and all of its vertices. Given a subgraph \mathbf{G} of a map \mathbf{M} , the *dual subgraph* of \mathbf{G} , denoted by \mathbf{G}^\dagger , is the subgraph of \mathbf{M}^\dagger consisting of all the edges that do not intersect \mathbf{G} . Following the terminology in [6], we call *subgraph-rooted map* a rooted planar map with a distinguished subgraph. Fig. 1(a) gives an example of a subgraph-rooted map with its dual map.

Local limit. For subgraph-rooted maps, the local distance is defined by

$$d_{\text{loc}}((\mathbf{M}, \mathbf{G}), (\mathbf{M}', \mathbf{G}')) = \inf \left\{ 2^{-R} \mid R \in \mathbb{N}, B_R(\mathbf{M}, \mathbf{G}) = B_R(\mathbf{M}', \mathbf{G}') \right\} \quad (1)$$

where $B_R(\mathbf{M}, \mathbf{G})$, the ball of radius R in (\mathbf{M}, \mathbf{G}) , is the subgraph-rooted map consisting of all vertices of \mathbf{M} at graph distance at most R from the root vertex and the edges between them. An edge of $B_R(\mathbf{M}, \mathbf{G})$ belongs to the distinguished subgraph of $B_R(\mathbf{M}, \mathbf{G})$ if and only if it is in \mathbf{G} . The space of all finite subgraph-rooted maps is not complete with respect to d_{loc} and we denote by \mathcal{M} its Cauchy completion. We call *infinite subgraph-rooted map* the elements of \mathcal{M} which are not finite subgraph-rooted map. Note that with this definition all infinite maps are *locally finite*, that is, every vertex is of finite degree.

The study of infinite random maps goes back to the works of Angel, Benjamini and Schramm on the Uniform Infinite Planar Triangulation (UIPT) [4, 2] obtained as the local limit of uniform

*IPhT, CEA-Saclay and Université Paris-Sud, E-mail: linxiao.chen@math.u-psud.fr




Figure 1: (a) A subgraph-rooted map and its dual. Edges of the distinguished subgraph are drawn in solid line, and the other edges in dashed line. The root corner is indicated by an arrow. (b) Loops separating the distinguished subgraph and its dual subgraph.

triangulations of size tending to infinity. Since then variants of this theorem have been proved for different classes of maps [13, 25, 26, 15, 8]. A common point of these infinite random lattices is that they are constructed from the uniform distribution on some finite sets of planar maps. In this work, we consider a different type of distribution.

cFK random map. For $n \geq 1$ we write \mathcal{M}_n for the set of all subgraph-rooted maps with n edges. Recall that in a dual subgraph-rooted maps, the distinguished subgraph \mathbf{G} and its dual subgraph \mathbf{G}^\dagger do not intersect. Therefore we can draw a set of loops tracing the boundary between them, as in Fig. 1(b). Let $\ell(\mathbf{M}, \mathbf{G})$ be the number of loops separating \mathbf{G} and \mathbf{G}^\dagger . For each $q > 0$, let $\mathbb{Q}_n^{(q)}$ be the probability distribution on \mathcal{M}_n defined by

$$\mathbb{Q}_n^{(q)}(\mathbf{M}, \mathbf{G}) \propto q^{\frac{1}{2}\ell(\mathbf{M}, \mathbf{G})} \quad (2)$$

By taking appropriate limits, we can define $\mathbb{Q}_n^{(q)}$ for $q \in \{0, \infty\}$. A *critical Fortuin-Kasteleyn (cFK) random map* of size n and of parameter q is a random variable of law $\mathbb{Q}_n^{(q)}$ (see Equation (5) below for the connection with the Fortuin-Kasteleyn random cluster model). From the definition of the loop number ℓ , it is easily seen that the law $\mathbb{Q}_n^{(q)}$ is self-dual (which is why we call it critical):

$$\mathbb{Q}_n^{(q)}(\mathbf{M}, \mathbf{G}) = \mathbb{Q}_n^{(q)}(\mathbf{M}^\dagger, \mathbf{G}^\dagger) \quad (3)$$

Our main result is:

Theorem 1. For each $q \in [0, \infty]$, we have $\mathbb{Q}_n^{(q)} \xrightarrow[n \rightarrow \infty]{} \mathbb{Q}_\infty^{(q)}$ in distribution with respect to the metric d_{loc} . Moreover, if (\mathbf{M}, \mathbf{G}) has law $\mathbb{Q}_\infty^{(q)}$, then

- $(\mathbf{M}, \mathbf{G}) = (\mathbf{M}^\dagger, \mathbf{G}^\dagger)$ in distribution,
- the map \mathbf{M} is almost surely one-ended and recurrent for the simple random walk,
- the laws of \mathbf{M} for different values of q are mutually singular.

So far two classes of methods have been developed to prove local convergence of finite random maps. The first one, initially used in [2] is based on precise asymptotic enumeration formulas for certain classes of maps. Although enumeration results about (a generalization of) cFK decorated maps have been obtained using combinatorial techniques [24, 17, 7, 19, 11, 10, 9], we are not going to follow this approach here. Instead, we will first transform our finite map model through a bijection into simpler objects. The archetype of such bijection is the famous Cori-Vauquelin-Schaeffer bijection and its generalizations [28, 12]. Then we take local limits of these simpler objects and construct the limit of the maps directly from the latter. This technique has been used e.g. in [13, 15, 8]. In this work the role of the Schaeffer bijection will be played by Sheffield's hamburger-cheeseburger bijection [29] which maps a cFK random map to a *random word* in a measure-preserving way. We will then construct the local limit of cFK random maps by showing that the random word converges locally to a limit, and that the hamburger-cheeseburger bijection has an almost surely continuous extension for that limit.

The cFK random maps have also been the subject of the recent works [5] and [21, 22, 23]. These works focused on finer properties of the infinite cFK random map such as exact scaling exponents or scaling limit of the model. In particular, the scaling exponents associated with the length and the enclosed area of a loop in the infinite cFK random map were derived independently. The main purpose of the present paper is to prove the local convergence of finite cFK maps to the infinite

cFK map. We offer a detailed proof and construct explicitly the infinite-volume version of the hamburger-cheeseburger bijection. The one-endedness and recurrence of the infinite cFK random map are obtained as a by-product of this bijection. The fact that the joint law of (\mathbf{M}, \mathbf{G}) is mutually singular for different q follows from the various scaling exponents computed in [5] and [21]. By replacing the law of (\mathbf{M}, \mathbf{G}) by its marginal in \mathbf{M} , we improve slightly the result. Our proof is based on an ergodicity result of the cFK random maps, which is of independent interest (See Appendix A).

The rest of this paper is organized as follows. In Section 2 we discuss the law of the cFK random map in more details and examine three interesting special cases. In Section 3 we first define the random word model underlying the hamburger-cheeseburger bijection. Then we show that the model has an explicit local limit, and we prove some properties of the limit. In Section 4 we construct the hamburger-cheeseburger bijection and prove Theorem 1 by translating the properties of the infinite random word in terms of the maps.

Acknowledgements. I deeply thank Jérémie Bouttier and Nicolas Curien for their many helps in completing this work. I am grateful to Ewain Gwynne for pointing out a mistake in the previous version of this paper. I thank the anonymous referees for many valuable comments, especially for suggesting the mutual singularity property stated in Theorem 1. I thank also the Isaac Newton Institute and the organizers of the Random Geometry programme for their hospitality during the completion of this work. This work is partly supported by Grant ANR-12-JS02-0001 (ANR CARTAPLUS) and by Grant ANR-14-CE25-0014 (ANR GRAAL).

2 More on cFK random map

Let (\mathbf{M}, \mathbf{G}) be a subgraph-rooted map and denote by $c(\mathbf{G})$ the number of connected components in \mathbf{G} . Recalling the definition of $\ell(\mathbf{M}, \mathbf{G})$ given in the introduction, it is not difficult to see that $\ell(\mathbf{M}, \mathbf{G}) = c(\mathbf{G}) + c(\mathbf{G}^\dagger) - 1$. However $c(\mathbf{G}^\dagger)$ is nothing but the number of faces of \mathbf{G} , therefore by Euler's relation we have

$$\ell(\mathbf{M}, \mathbf{G}) = e(\mathbf{G}) + 2c(\mathbf{G}) - v(\mathbf{M}), \quad (4)$$

where $e(\mathbf{G})$ is the number of edges in \mathbf{G} , and $v(\mathbf{M})$ is the number of vertices in \mathbf{M} . This gives the following expression of the first marginal of $\mathbb{Q}_n^{(q)}$: for all rooted map \mathbf{M} with n edges, we have

$$\mathbb{Q}_n^{(q)}(\mathbf{M}) \propto q^{-\frac{1}{2}v(\mathbf{M})} \sum_{\mathbf{G} \subset \mathbf{M}} \sqrt{q}^{e(\mathbf{G})} q^{c(\mathbf{G})}. \quad (5)$$

The sum on the right-hand side over all the subgraphs of \mathbf{M} is precisely the partition function of the Fortuin-Kasteleyn random cluster model or, equivalently, of the Potts model on the map \mathbf{M} (The two partition functions are equal. See e.g. [18, Section 1.4]. See also [9, Section 2.1] for a review of their connection with loop models on planar lattices). For this reason, the cFK random map is used as a model of quantum gravity theory in which the geometry of the space interacts with the matter (spins in the Potts model). Note that the “temperature” in the Potts model and the prefactor $q^{-\frac{1}{2}v(\mathbf{M})}$ in (5) are tuned to ensure self-duality, which is crucial for our result to hold.

Three values of the parameter q deserve special attention, since the cFK random map has nice combinatorial interpretations in these cases.

$q = 0$: $\mathbb{Q}_n^{(0)}$ is the uniform measure on the elements of \mathcal{M}_n which minimize the number of loops ℓ . The minimum is $\ell_{\min} = 1$ and it is achieved if and only if the subgraph \mathbf{G} is a spanning tree of \mathbf{M} . Therefore under $\mathbb{Q}_n^{(0)}$, the map \mathbf{M} is chosen with probability proportional to the number of its spanning trees, and conditionally on \mathbf{M} , \mathbf{G} is a uniform spanning tree of \mathbf{M} .

At the limit, the marginal law of \mathbf{G} under $\mathbb{Q}_\infty^{(0)}$ will be that of a critical geometric Galton-Walton tree conditioned to survive. This will be clear once we defined the hamburger-cheeseburger bijection. In fact when $q = 0$, the hamburger-cheeseburger bijection is reduced to a bijection between tree-rooted maps and excursions of simple random walk on \mathbb{Z}^2 introduced earlier by Bernardi [6].

$q = 1$: $\mathbb{Q}_n^{(1)}$ is the uniform measure on \mathcal{M}_n . Since each planar map with n edges has 2^n subgraphs, \mathbf{M} is a uniform planar map chosen among the maps with n edges. Thus in the case $q = 1$, Theorem 1 can be seen as a construction of the Uniform Infinite Planar Map or of the Uniform Infinite Planar Quadrangulation via Tutte's bijection. It is a curious fact that with this approach, one has to first decorate a uniform planar map with a random subgraph in order to show the local convergence of the map. As we will see later, the couple (\mathbf{M}, \mathbf{G}) is encoded by the hamburger-cheeseburger bijection in an entangled way.

$q = \infty$: Similarly to the case $q = 0$, the probability $\mathbb{Q}_n^{(\infty)}$ is the uniform measure on the elements of \mathcal{M}_n which maximize ℓ . To see what are these elements, remark that each connected component of \mathbf{G} contains at least one vertex, therefore

$$c(\mathbf{G}) \leq v(\mathbf{M}) \quad (6)$$

And, at least one edge must be removed from \mathbf{M} to create a new connected component, so

$$c(\mathbf{G}) \leq c(\mathbf{M}) + e(\mathbf{M}) - e(\mathbf{G}) = n + 1 - e(\mathbf{G}) \quad (7)$$

Summing the two relations, we see that the maximal number of loops is $\ell_{\max} = n + 1$ and it is achieved if and only if each connected component of \mathbf{G} contains exactly one vertex (i.e. all edges of \mathbf{G} are self-loops) *and* that the complementary subgraph $\mathbf{M} \setminus \mathbf{G}$ is a tree. Fig. 2(a) gives an example of such couple (\mathbf{M}, \mathbf{G}) .




Figure 2: (a) A subgraph-rooted map which maximizes the number of loops ℓ . Colors are used only to illustrate the bijection. (b) The percolation configuration on a rooted tree associated to this map by the bijection. The divided vertices, as well as the replaced edges, are drawn in the same color before and after the bijection.

This model of loop-decorated tree is in bijection with bond percolation of parameter $1/2$ on a uniform random plane tree with n edges, as we now explain. For a couple (\mathbf{M}, \mathbf{G}) satisfying the above conditions, consider a self-loop \mathbf{e} in \mathbf{G} . This self-loop separates the rest of the map \mathbf{M} into two parts which share only the vertex of \mathbf{e} . We divide this vertex in two, and replace the self-loop \mathbf{e} by an edge joining the two child vertices. The new edge is always considered part of \mathbf{G} . By repeating this operation for all self-loops in the subgraph \mathbf{G} in an arbitrary order, we transform the map \mathbf{M} into a rooted plane tree, see Fig. 2. This gives a bijection from the support of $\mathbb{Q}_n^{(\infty)}$ to the set of rooted plane tree of n edges with a distinguished subgraph. The latter object converges locally to a critical geometric Galton-Watson tree conditioned to survive, in which each edge belongs to the distinguished subgraph with probability $1/2$ independently from other edges. Using the inverse of the bijection above (which is almost surely continuous at the limit), we can explicit the law $\mathbb{Q}_\infty^{(\infty)}$. In particular, it is easily seen that \mathbf{M} is almost surely a one-ended tree plus finitely many self-loops at each vertex. Therefore it is one-ended and recurrent.

3 Local limit of random words

In this section we define the random word model underlying the hamburger-cheeseburger bijection, and establish its local limit.

We consider words on the alphabet $\Theta = \{\mathbf{a}, \mathbf{b}, \mathbf{A}, \mathbf{B}, \mathbf{F}\}$. Formally, a *word* w is a mapping from an interval I of integers to Θ . We write $w \in \Theta^I$ and we call I the *domain* of w . Let \mathcal{W} be the space of all words, that is,

$$\mathcal{W} = \bigcup_I \Theta^I \quad (8)$$

where I runs over all subintervals of \mathbb{Z} . Note that a word can be finite, semi-infinite or bi-infinite. We denote by \emptyset the *empty word*. Given a word w of domain I and $k \in I$, we denote by w_k the letter of index k in w . More generally, if J is an (integer or real) interval, we denote by w_J the *restriction* of the word w to $I \cap J$. For example, if $w = \mathbf{bAbafABa} \in \Theta^{\{0, \dots, 7\}}$, then $w_{[2,6)} = \mathbf{baFA} \in \Theta^{\{2,3,4,5\}}$. We endow \mathcal{W} with the local distance

$$D_{\text{loc}}(w, w') = \inf \left\{ 2^{-R} \mid R \in \mathbb{N}, w_{[-R,R)} = w'_{[-R,R)} \right\} \quad (9)$$

Note that the equality $w_{[-R,R)} = w'_{[-R,R)}$ implies that $I \cap [-R, R) = I' \cap [-R, R)$, where I (resp. I') is the domain of the word w (resp. w'). It is easily seen that $(\mathcal{W}, D_{\text{loc}})$ is a compact metric space.

3.1 Reduction of words

Now we define the reduction operation on the words. For each word w , this operation specifies a pairing between letters in the word called *matching*, and returns two shorter words \bar{w}^λ and \bar{w}^Λ .

We follow the exposition given in [29]. The letters $\mathbf{a}, \mathbf{b}, \mathbf{A}, \mathbf{B}, \mathbf{F}$ are interpreted as, respectively, a hamburger, a cheeseburger, a hamburger order, a cheeseburger order and a *flexible* order. They obey the following order fulfillment relation: a hamburger order \mathbf{A} can only be fulfilled by a hamburger \mathbf{a} , a cheeseburger order \mathbf{B} by a cheeseburger \mathbf{b} , while a flexible order \mathbf{F} can be fulfilled either by a hamburger \mathbf{a} or by a cheeseburger \mathbf{b} . We write $\lambda = \{\mathbf{a}, \mathbf{b}\}$ and $\Lambda = \{\mathbf{A}, \mathbf{B}, \mathbf{F}\}$ for the set of lowercase letters (burgers) and uppercase letters (orders).

Finite case. A finite word $w \in \Theta^I$ can be seen from left to right as a sequence of events that happen in a restaurant with time indexed by I . Namely, at each time $k \in I$, either a burger is produced, or an order is placed. The restaurant puts all its burgers on a stack S , and takes note of unfulfilled orders in a list L . Both S and L start as the empty string. When a burger is produced, it is appended at the end of the stack. When an order arrives, we check if it can be fulfilled by one of the burgers in the stack. If so, we take the *last* such burger in the stack and fulfills the order. (That is, the stack is last-in-first-out.) Otherwise, the order goes to the end of the list L . Fig. 3 illustrates this dynamics with an example.




Figure 3: The reduction procedure of a word and the associated arch diagram

We encode the *matching* of w by a function $\phi_w : I \rightarrow I \cup \{-\infty, \infty\}$. If the burger produced at time j is consumed by an order placed at time k , then the letters w_j and w_k are said to be *matched*, and we set $\phi_w(j) = k$ and $\phi_w(k) = j$. On the other hand, if a letter w_k corresponds to a unfulfilled order or a leftover burger, then it is *unmatched*, and we set $\phi_w(k) = \infty$ if it is a burger ($w_k \in \lambda$) and $\phi_w(k) = -\infty$ if it is an order ($w_k \in \Lambda$).

Moreover, let us denote by \bar{w}^Λ (resp. \bar{w}^λ) the state of the list L (resp. the stack S) at the end of the day. Together they give the reduced form of the word w .

Definition 2 (reduced word). The reduced word associated to a finite word w is the concatenation $\bar{w} = \bar{w}^\Lambda \bar{w}^\lambda$. That is, it is the list of unmatched uppercase letters in w , followed by the list of unmatched lowercase letters in w .

The matching and the reduced word can be represented as an arch diagram as follows. For each letter w_j in the word w , draw a vertex in the complex plane at position j . For each pair of matched letters w_j and w_k , draw a semi-circular arch that links the corresponding pair of vertices. This arch is drawn in the upper half plane if it is incident to an \mathbf{a} -vertex, and in the lower half plane if it is incident to a \mathbf{b} -vertex. For an unmatched letter w_j , we draw an open arch from j tending to the left if $\phi_w(j) = -\infty$, or to the right if $\phi_w(j) = \infty$. See Fig. 3.

It should be clear from the definition of matching operation that the arches in this diagram do not intersect each other. We shall come back to this diagram in Section 4 to construct the hamburger-cheeseburger bijection.

Infinite case. Remark that a hamburger produced at time j is consumed by a hamburger order at time $k > j$ if and only if 1) all the hamburgers produced during the interval $[j+1, k-1]$ are consumed strictly before time k , and 2) all the hamburger or flexible orders placed during $[j+1, k-1]$ are fulfilled by a burger produced strictly after time j . In terms of the reduced word, this means that two letters $w_j = \mathbf{a}$ and $w_k = \mathbf{A}$ are matched if and only if $\bar{w}_{(j,k)}$ does not contain any \mathbf{a} , \mathbf{A} or \mathbf{F} . This can be generalized to any pair of burger/order.

Proposition 3 ([29]). For $j < k$, assume that $w_j \in \lambda$ and $w_k \in \Lambda$ can be matched. Then they are matched in w if and only if $\bar{w}_{(j,k)}$ does not contain any letter that can be matched to either w_j or w_k .

This shows that the matching rule is entirely determined by the reduction operator. More importantly, we see that the matching rule is *local*, that is, whether $\phi_w(j) = k$ or not only depends on $w_{[j,k]}$. From this we deduce that the reduction operator is compatible with string concatenation, that is, $\overline{uv} = \bar{u}\bar{v} = \bar{u}v$ for any pair of finite words u, v .

This locality property allows us to define ϕ_w for infinite words w . Then, we can also read \bar{w}^λ (resp. \bar{w}^Λ) from ϕ_w as the (possibly infinite) sequence of unmatched lowercase (resp. uppercase) letters. However \bar{w} is not defined in general, since the concatenation $\bar{w}^\Lambda \bar{w}^\lambda$ does not always make sense.

Random word model and local limit. For each $p \in [0, 1]$, let $\theta^{(p)}$ be the probability measure on Θ such that

$$\theta^{(p)}(\mathbf{a}) = \theta^{(p)}(\mathbf{b}) = \frac{1}{4} \quad \theta^{(p)}(\mathbf{A}) = \theta^{(p)}(\mathbf{B}) = \frac{1-p}{4} \quad \theta^{(p)}(\mathbf{F}) = \frac{p}{2}$$

Here p should be interpreted as the proportion of flexible orders among all the orders. Remark that, regardless of the value of p , the distribution is symmetric when exchanging \mathbf{a} with \mathbf{b} and \mathbf{A} with \mathbf{B} . As we will see in Section 4, this corresponds to the self-duality of cFK random maps.

For $n \geq 1$, let $I_k = \{-k, \dots, 2n-1-k\}$, and set

$$\mathcal{W}_n = \bigcup_{0 \leq k < 2n} \left\{ w \in \Theta^{I_k} \mid \overline{w} = \emptyset \right\} \quad (10)$$

For $p \in [0, 1]$, let $\mathbb{P}_n^{(p)}$ be the probability measure on \mathcal{W}_n proportional to the direct product of $\theta^{(p)}$, that is, for all $w \in \mathcal{W}_n$,

$$\mathbb{P}_n^{(p)}(w) \propto \prod_j \theta^{(p)}(w_j) \quad (11)$$

where the product is taken over the domain of w . In addition, let $\mathbb{P}_\infty^{(p)} = \theta^{(p)} \otimes \mathbb{Z}$ be the product measure on bi-infinite words. Our proof of Theorem 1 relies mainly on the following proposition, stated by Sheffield in an informal way in [29].

Proposition 4. *For all $p \in [0, 1]$, we have $\mathbb{P}_n^{(p)} \rightarrow \mathbb{P}_\infty^{(p)}$ in law for D_{loc} as $n \rightarrow \infty$.*

3.2 Proof of Proposition 4

We follow the approach proposed by Sheffield in [29, Section 4.2]. Let $W^{(p)}$ be a random word of law $\mathbb{P}_\infty^{(p)}$, so that $W_{[0,n]}^{(p)}$ is a word of length n with i.i.d. letters. By compactness of $(\mathcal{W}, D_{\text{loc}})$, it suffices to show that for any ball B_{loc} in this space, we have $\mathbb{P}_n^{(p)}(B_{\text{loc}}) \rightarrow \mathbb{P}_\infty^{(p)}(B_{\text{loc}})$. Note that D_{loc} is an ultrametric and the ball $B_{\text{loc}}(w, 2^{-R})$ of radius 2^{-R} around w is the set of words which are identical to w when restricted to $[-R, R]$. In the rest of the proof, we fix an integer $R \geq 1$ and a word $w \in \Theta^{[-R,R] \cap \mathbb{Z}}$. Recall that $W^{(p)}$ has law $\mathbb{P}_\infty^{(p)}$. In the following we omit the parameter p from the superscripts to keep simple notations.

Recall that the space \mathcal{W}_n is made up of $2n$ copies of the set $\{w \in \Theta^{2n} \mid \overline{w} = \emptyset\}$ differing from each other by translation of the indices. Therefore \mathbb{P}_n can be seen as the conditional law of $W_{[-K, 2n-K]}$ on the event $\{\overline{W}_{[-K, 2n-K]} = \emptyset\}$, where K is a uniform random variable on $\{0, \dots, 2n-1\}$ independent from W . Moreover, for the word $W_{[-K, 2n-K]}$ to have w as its restriction to $[-R, R]$, one must to have $R \leq K \leq 2n-R$. Hence,

$$\begin{aligned} \mathbb{P}_n \left(B_{\text{loc}}(w, 2^{-R}) \right) &= \mathbb{P} \left(R \leq K \leq 2n-R \text{ and } W_{[-R,R]} = w \mid \overline{W}_{[-K, 2n-K]} = \emptyset \right) \\ &= \frac{1}{2n} \sum_{k=R}^{2n-R} \mathbb{P} \left(W_{[-R,R]} = w \mid \overline{W}_{[-k, 2n-k]} = \emptyset \right) \\ &= \frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{P} \left(W_{[k, k+2R]} \simeq w \mid \overline{W}_{[0, 2n]} = \emptyset \right) \\ &= \mathbb{E} \left[\frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{1}_{\{W_{[k, k+2R]} \simeq w\}} \mid \overline{W}_{[0, 2n]} = \emptyset \right] \end{aligned}$$

where in the last two steps, we denote by $u \simeq v$ the fact that two words are equal up to an overall translation of indices. On the other hand, set

$$\pi_w = \mathbb{P}(B_{\text{loc}}(w, 2^{-R})) = \prod_{k=-R}^{R-1} \theta(w_k) \quad (12)$$

By translation invariance of W we have

$$\mathbb{E} \left[\frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{1}_{\{W_{[k, k+2R]} \simeq w\}} \right] = \frac{2n-2R+1}{2n} \pi_w \quad (13)$$

In fact, up to boundary terms of the order $O(R/n)$, the quantity inside the expectation is the empirical measure of the Markov chain $(W_{[k, k+2R]})_{k \geq 0}$ taken at the state w . This is an irreducible Markov chain on the finite state space Θ^{2R} . Sanov's theorem (see e.g. [16, Theorem 3.1.2]) gives the following large deviation estimate. For any $\epsilon > 0$, there are constants $A_\epsilon, C_\epsilon > 0$ depending only on ϵ and on the transition matrix of $(W_{[k, k+2R]})_{k \geq 0}$, such that

$$\mathbb{P} \left(\left| \frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{1}_{\{W_{[k, k+2R]} \simeq w\}} - \pi_w \right| > \epsilon \right) \leq A_\epsilon e^{-C_\epsilon n} \quad (14)$$

for all $n \geq 1$. Since $\left| \frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{1}_{\{W_{[k, k+2R]} \simeq w\}} - \pi_w \right|$ is bounded by 1, we have

$$\begin{aligned} \left| \mathbb{P}_n \left(B_{\text{loc}}(w, 2^{-R}) \right) - \pi_w \right| &\leq \mathbb{E} \left[\left| \frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{1}_{\{W_{[k, k+2R]} \simeq w\}} - \pi_w \right| \middle| \overline{W}_{[0, 2n]} = \emptyset \right] \\ &\leq \epsilon + \mathbb{P} \left(\left| \frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{1}_{\{W_{[k, k+2R]} \simeq w\}} - \pi_w \right| > \epsilon \middle| \overline{W}_{[0, 2n]} = \emptyset \right) \\ &\leq \epsilon + \frac{1}{\mathbb{P}(\overline{W}_{[0, 2n]} = \emptyset)} \cdot \mathbb{P} \left(\left| \frac{1}{2n} \sum_{k=0}^{2n-2R} \mathbb{1}_{\{W_{[k, k+2R]} \simeq w\}} - \pi_w \right| > \epsilon \right) \\ &\leq \epsilon + \frac{A_\epsilon e^{-C_\epsilon n}}{\mathbb{P}(\overline{W}_{[0, 2n]} = \emptyset)} \end{aligned}$$

According to [29, Eq. (28)],¹

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\overline{W}_{[0, 2n]} = \emptyset) = 0$$

Therefore the second term converges to zero as $n \rightarrow \infty$. Since ϵ can be taken arbitrarily close to zero, this shows that $\mathbb{P}_n(B_{\text{loc}}(w, 2^{-R})) \rightarrow \pi_w$ as $n \rightarrow \infty$.

3.3 Some properties of the limiting random word

In this section we show two properties of the infinite random word $W^{(p)}$ which will be the word-counterpart of Theorem 1. Both properties are true for general $p \in [0, 1]$. However we will only write proofs for $p < 1$, since the case $p = 1$ corresponds to cFK random maps with parameter $q = \infty$, for which the local limit is explicit. (The proofs for $p = 1$ are actually easier, but they require different arguments.)

Proposition 5 (Sheffield [29]). *For all $p \in [0, 1]$, almost surely,*

1. $\overline{W^{(p)}} = \emptyset$, that is, every letter in $W^{(p)}$ is matched.
2. For all $k \in \mathbb{Z}$, $\overline{W_{(-\infty, k]}^{(p)}}$ contains infinitely many **a** and infinitely many **b**.

Proof. The first assertion is proved as Proposition 2.2 in [29]. For the second assertion, recall that $\overline{W_{(-\infty, k]}^{(p)}}$ represents a left-infinite stack of burgers. Now assume for some $k \in \mathbb{Z}$, it contains only N letters **a** with positive probability. Then, with probability $(\frac{1-p}{4})^{N+1}$ and independently of $W_{(-\infty, k]}^{(p)}$, all the $N+1$ letters in $W_{[k, k+N]}^{(p)}$ are **A**. This will leave the **A** at position $k+N$ unmatched in W , which happens with zero probability according to the first assertion. This gives a contradiction when $p < 1$. \square

For each random word $W^{(p)}$, consider a random walk Z on \mathbb{Z}^2 starting from the origin: $Z_0 = (0, 0)$, and for all $k \in \mathbb{Z}$,

$$Z_{k+1} - Z_k = \begin{cases} (1, 0) & \text{if } W_k^{(p)} = \mathbf{a} \\ (-1, 0) & \text{if } W_k^{(p)} \text{ is matched to an } \mathbf{a} \\ (0, 1) & \text{if } W_k^{(p)} = \mathbf{b} \\ (0, -1) & \text{if } W_k^{(p)} \text{ is matched to a } \mathbf{b} \end{cases} \quad (15)$$

By Proposition 5, Z_k is almost surely well-defined for all $k \in \mathbb{Z}$. A lot of information about the random word W can be read from Z . The main result of [29] shows that under diffusive rescaling, Z converges to a Brownian motion in \mathbb{R}^2 with a diffusivity matrix that depends on p , demonstrating a phase transition at $p = 1/2$.

Let $(X, Y) = Z$. Then X represents the net hamburger count and Y the net cheeseburger count. Set $i_0 = \sup \{i < 0 \mid X_i = -1\}$ and $j_0 = \inf \{j > 0 \mid X_j = -1\}$. Let N_0 be the number of times that X visits the state 0 between time i_0 and j_0 . We shall see in Section 4.2 that N_0 is exactly the degree of root vertex in the infinite cFK-random map. Below we prove that the distribution of N_0 has an exponential tail, that is, there exists constants A and $c > 0$ such that $\mathbb{P}(N_0 \geq x) \leq A e^{-cx}$ for all $x \geq 0$.

Proposition 6. *N_0 has an exponential tail distribution for all $p \in [0, 1]$.*

Proof. First let us consider N_0^+ , the number of times that X visits the state 0 between time 0 and j_0 . Remark that at positive time, the process X is adapted to the natural filtration $(\mathcal{F}_k)_{k \geq 0}$, where

¹It has been shown in [21] that $\mathbb{P}(\overline{W}_{[0, 2n]} = \emptyset)$ decays as a power of n , with the exact exponent as a function of p . But we do not need this fact here.




Figure 4: The decomposition of N_0^+ into intervals $[S_k, T_k]$

\mathcal{F}_k is the σ -algebra generated by $W_{[0,k]}$. Define two sequences of \mathcal{F} -stopping times $(S_m)_{m \geq 0}$ and $(T_m)_{m \geq 0}$ by $S_0 = 0$, and that for all $m \geq 0$,

$$\begin{aligned} T_m &= \inf \{k > S_m \mid X_k \neq 0\} \\ S_{m+1} &= \inf \{k > T_m \mid X_k = 0\} \end{aligned}$$

The sequence S (resp. T) marks the times that X arrives at (resp. departs from) the state 0. Therefore the total number of visits of the state 0 between time 0 and T_m is $\sum_{i=0}^m (T_i - S_i)$, see Fig. 4.

By construction, j_0 is the smallest T_m such that $X_{T_m} = -1$. On the other hand, we have $X_{T_m} \in \{-1, +1\}$ for all m and

$$\begin{aligned} X_{T_m} = +1 &\Leftrightarrow W_{T_m} = \mathbf{a} \\ X_{T_m} = -1 &\Leftrightarrow W_{T_m} = \mathbf{A} \text{ or } W_{T_m} \text{ is an } \mathbf{F} \text{ matched to an } \mathbf{a} \end{aligned}$$

Consider the stopping time

$$M = \inf \{m \geq 0 \mid W_{T_m} = \mathbf{A}\} \quad (16)$$

Then we have $j_0 \leq T_M$, and therefore

$$N_0^+ \leq \sum_{m=0}^M (T_m - S_m) \quad (17)$$

On the other hand, M is the smallest m such that, starting from time S_m , an \mathbf{A} comes before an \mathbf{a} . Therefore M is a geometric random variable of mean $\frac{1-p}{p}$.

Assume $p < 1$ so that M is almost surely finite. Fix an integer $m \geq 0$. By the strong Markov property, conditionally to $\{M = m\}$, the sequence $(T_i - S_i, i = 0 \dots m-1)$ is i.i.d., and each term in the sequence has the same law as the first arrival time of \mathbf{a} in the sequence $(W_k)_{k \geq 0}$ conditioned not to contain \mathbf{A} . In other words, conditionally to $\{M = m\}$, $(T_i - S_i, i = 0 \dots m-1)$ is an i.i.d. sequence of geometric random variables of mean $\frac{1}{2+p}$. Similarly, conditionally to $\{M = m\}$, $T_m - S_m$ is a geometric random variable of mean $\frac{1-p}{2+p}$ independent from the sequence $(T_i - S_i, i = 0 \dots m-1)$.

Then, a direct computation shows that the exponential moment $\mathbb{E}[e^{\lambda N_0^+}]$ is finite for some $\gamma > 0$. And by Markov's inequality, the distribution of N_0^+ has an exponential tail when $p < 1$.

Now we claim that conditionally to the value of N_0 , the variable N_0^+ is uniform on $\{1, \dots, N_0\}$ which implies that N_0 also has an exponential tail distribution. To see why the conditional law is uniform, consider N_0 and N_0^+ for finite words defined in the same way as for the infinite word $W_\infty^{(p)}$. Note that for a finite word w the process X does not necessarily hit -1 at negative (resp. positive) times. In this case we just replace i_0 (resp. j_0) by the infimum (resp. supremum) of the domain of w . Then, $w \mapsto (N_0, N_0^+)$ is a $D_{1\text{oc}}$ -continuous function defined on the union of $\cup_{n \geq 0} \mathcal{W}_n$ and the support of $W_\infty^{(p)}$. Therefore for any integers $k \leq m$,

$$\mathbb{P}_n^{(p)}(N_0 = m, N_0^+ = k) \xrightarrow{n \rightarrow \infty} \mathbb{P}_\infty^{(p)}(N_0 = m, N_0^+ = k) \quad (18)$$

But, given the sequence of letters in a word, the law $\mathbb{P}_n^{(p)}$ chooses the letter of index 0 uniformly at random among all the letters. A simple counting shows that for all $1 \leq k, k' \leq m$, we have $\mathbb{P}_n^{(p)}(N_0 = m, N_0^+ = k) = \mathbb{P}_n^{(p)}(N_0 = m, N_0^+ = k')$. Letting $n \rightarrow \infty$ shows that the conditional law of N_0^+ given N_0 is uniform under $\mathbb{P}_\infty^{(p)}$. \square

4 The hamburger-cheeseburger bijection

4.1 Construction

In this section we present (a slight variant of) the hamburger-cheeseburger bijection of Sheffield. We refer to [29] for the proof of bijectivity and for historical notes.

We define the hamburger-cheeseburger bijection $\tilde{\Psi}$ on a subset of the space \mathcal{W} , and it takes values in the space $\tilde{\mathcal{M}}$ of *doubly-rooted planar maps with a distinguished subgraph*, that is, planar maps with two distinguished corners and one distinguished subgraph. We can write this space as

$$\tilde{\mathcal{M}} = \{(\mathbf{M}, \mathbf{G}, \mathbf{s}) \mid (\mathbf{M}, \mathbf{G}) \in \mathcal{M} \text{ and } \mathbf{s} \text{ is a corner of } \mathbf{M}\} \quad (19)$$

Note that the second root \mathbf{s} may be equal to or different from the root of \mathbf{M} . We define in the same way $\tilde{\mathcal{M}}_n$, the doubly-rooted version of the space \mathcal{M}_n . Its cardinal is $2n$ times that of \mathcal{M}_n .

We start by constructing $\tilde{\Psi} : \mathcal{W}_n \rightarrow \tilde{\mathcal{M}}_n$ in three steps. The first step transforms a word in \mathcal{W}_n into a decorated planar map called *arch graph*. The second and the third step apply graph duality and local transformations to the arch graph to get a *tree-rooted map*, and then a *subgraph-rooted map* in $\tilde{\mathcal{M}}_n$.

Step 1: from words to arch graphs. Fix a word $w \in \mathcal{W}_n$. Recall from Section 3 the construction of the non-crossing arch diagram associated to w . In particular since $\bar{w} = \emptyset$, there is no half-arch. We link neighboring vertices by unit segments $[j-1, j]$ and link the first vertex to the last vertex by an edge that wires around the whole picture on either side of the real axis, without intersecting any other edges. This defines a planar map \mathbf{A} of $2n$ vertices and $2n$ edges. In \mathbf{A} we distinguish edges coming from arches and the other edges. The latter forms a simple loop passing through all the vertices.

We further decorate \mathbf{A} with additional pieces of information. Recall that the word w is indexed by an interval of the form $I_k = \{-k, \dots, 2n-1-k\}$ where $0 \leq k < 2n$. We will mark the oriented edge \mathbf{r} from the vertex 0 to the vertex -1 , and the oriented edge \mathbf{s} from the first vertex $(-k)$ to the last vertex $(2n-1-k)$. If $k=0$, then \mathbf{r} and \mathbf{s} coincide. Furthermore, we mark each arch incident to an F-vertex by a star $*$. (See Fig. 5) We call the decorated planar map \mathbf{A} the *arch graph* of w . One can check that it completely determines the underlying word w .

Step 2: from arch graphs to tree-rooted maps. We now consider the dual map Δ of the arch graph \mathbf{A} . Let \mathbf{Q} be the subgraph of Δ consisting of edges whose dual edge is on the loop in \mathbf{A} . We denote by $\Delta \setminus \mathbf{Q}$ the set of remaining edges of Δ (that is, the edges intersecting one of the arches).

Proposition 7. *The map Δ is a triangulation, the map \mathbf{Q} is a quadrangulation with n faces and $\Delta \setminus \mathbf{Q}$ consists of two trees.* \square

We denote by \mathbf{T} and \mathbf{T}^\dagger the two trees in $\Delta \setminus \mathbf{Q}$, with \mathbf{T} corresponding to faces of the arch graph in the upper half plane. Then \mathbf{Q} , \mathbf{T} and \mathbf{T}^\dagger form a partition of edges in the triangulation Δ . Note that \mathbf{T} and \mathbf{T}^\dagger give the (unique) bipartition of vertices of \mathbf{Q} . Let \mathbf{M} be the planar map associated to \mathbf{Q} by Tutte's bijection, such that \mathbf{M} has the same vertex set as \mathbf{T} . (The latter prescription allows us to bypass the root, and define Tutte's bijection from unrooted quadrangulations to unrooted maps.) We thus obtain a couple (\mathbf{M}, \mathbf{T}) in which \mathbf{M} is a map with n edges and \mathbf{T} is spanning tree of \mathbf{M} . Remark that \mathbf{T}^\dagger is the dual spanning tree of \mathbf{T} in the dual map \mathbf{M}^\dagger . This relates the duality of maps with the duality on words which consists of exchanging \mathbf{a} with \mathbf{b} and \mathbf{A} with \mathbf{B} .

Fig. 5(a) summarizes the mapping from words to tree-rooted maps (Step 1 and 2) with an example. Note that we have omitted the two roots and the stars on the arch graph in the above discussion. But since graph duality and Tutte's bijection provide canonical bijections between edges, the roots and stars can be simply transferred from the arches in \mathbf{A} to the edges in \mathbf{M} . With the roots and stars taken into account, it is clear that $w \mapsto (\mathbf{M}, \mathbf{T})$ is a bijection from \mathcal{W}_n onto its image.

Step 3: from tree-rooted maps to subgraph-rooted maps. Now we “switch the status” of every starred edge in \mathbf{M} relative to the spanning tree \mathbf{T} . That is, if a starred edge is not in \mathbf{T} , we add it to \mathbf{T} ; if it is already in \mathbf{T} , we remove it from \mathbf{T} . Let \mathbf{G} be the resulting subgraph. See Fig. 5(b) for an example.

Recall that there are two marked corners \mathbf{r} and \mathbf{s} in the map \mathbf{M} . By an abuse of notation, from now on we denote by \mathbf{M} the *rooted* map with root corner \mathbf{r} . Then, the hamburger-cheeseburger bijection is defined by $\tilde{\Psi}(w) = (\mathbf{M}, \mathbf{G}, \mathbf{s})$. Let $\Psi(w) = (\mathbf{M}, \mathbf{G})$ be its projection obtained by forgetting the second root corner. We denote by $(\#F)w$ the number of letters F in w , and by ℓ the number of loops associated to the corresponding subgraph-rooted map (\mathbf{M}, \mathbf{G}) .

Theorem 8 (Sheffield [29]). *The mapping $\tilde{\Psi} : \mathcal{W}_n \rightarrow \tilde{\mathcal{M}}_n$ is a bijection such that $\ell = 1 + (\#F)w$ for all $w \in \mathcal{W}_n$. And $\mathbb{Q}_n^{(q)}$ is the image measure of $\mathbb{P}_n^{(p)}$ by Ψ whenever*

$$p = \frac{\sqrt{q}}{2 + \sqrt{q}}.$$

Proof. The proof of this can be found in [29]. However we include a proof of the second fact to enlighten the relation $p = \frac{\sqrt{p}}{2 + \sqrt{q}}$. For $w \in \mathcal{W}_n$, since $\bar{w} = \emptyset$, we have $(\#a)w + (\#b)w = (\#A)w +$

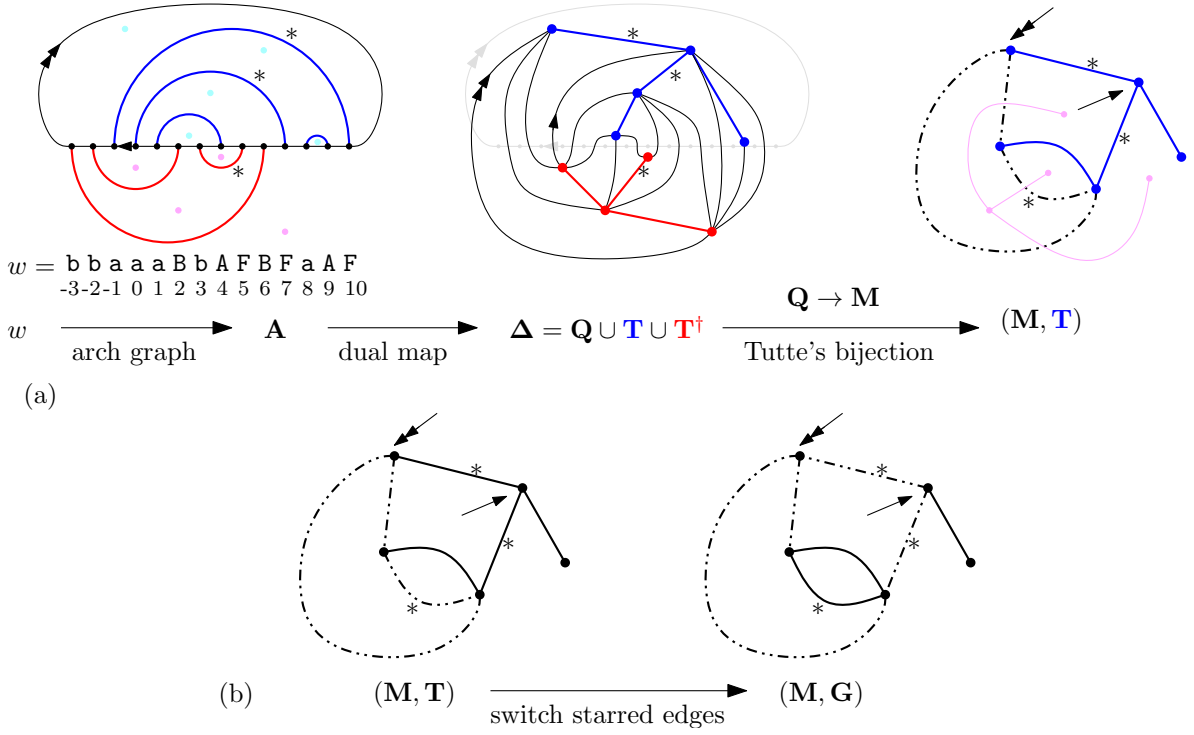


Figure 5: Construction of the hamburger-cheesburger bijection. (a) From word to tree-rooted map. (b) From tree-rooted map to subgraph-rooted map.

$(\#\mathbf{B})w + (\#\mathbf{F})w = n$. Therefore, when $p = \frac{\sqrt{q}}{2+\sqrt{q}}$,

$$\begin{aligned} \mathbb{P}_n^{(p)}(w) &\propto \left(\frac{1}{4}\right)^{(\#\mathbf{a})w + (\#\mathbf{b})w} \left(\frac{1-p}{4}\right)^{(\#\mathbf{A})w + (\#\mathbf{B})w} \left(\frac{p}{2}\right)^{(\#\mathbf{F})w} \\ &= \left(\frac{1}{4}\right)^n \left(\frac{1-p}{4}\right)^{n - (\#\mathbf{F})w} \left(\frac{p}{2}\right)^{(\#\mathbf{F})w} \propto \left(\frac{2p}{1-p}\right)^{(\#\mathbf{F})w} = \sqrt{q}^{\ell-1} \end{aligned}$$

After normalization, this shows that $\mathbb{Q}_n^{(q)}$ is the image measure of $\mathbb{P}_n^{(p)}$ by Ψ . \square

Proposition 9. *We can extend the mapping Ψ to $\mathcal{W} \rightarrow \mathcal{M}$ so that it is $\mathbb{P}_\infty^{(p)}$ -almost surely continuous with respect to D_{loc} and d_{loc} , for all $p \in [0, 1]$.*

Proof. Observe that if we do not care about the location of the second root \mathbf{s} , then the word w used in the construction of Ψ does not have to be finite. Set

$$\mathcal{W}_\infty = \left\{ w \in \Theta^\mathbb{Z} \mid \begin{array}{l} \bar{w} = \emptyset \text{ and for all } k \in \mathbb{Z}, \bar{w}_{(-\infty, k)}^\lambda \text{ contains} \\ \text{infinitely many } \mathbf{a} \text{ and infinitely many } \mathbf{b} \end{array} \right\} \quad (20)$$

We claim that indeed, for each $w \in \mathcal{W}_\infty$, Step 1, 2 and 3 of the construction define a (locally finite) infinite subgraph-rooted map: as in the case of finite words, the condition $\bar{w} = \emptyset$ ensures that the arch graph \mathbf{A} of w is a well-defined infinite planar map (that is, all the arches are closed). To see that its dual map Δ is a locally finite, infinite triangulation, we only need to check that each face of \mathbf{A} has finite degree. Observe that a letter \mathbf{a} in w appears in $\bar{w}_{(-\infty, k)}^\lambda$ if and only if it is on the left of w_k , and that its partner is on the right of w_k . This corresponds to an arch passing above the vertex k . Therefore, the remaining condition in the definition of \mathcal{W}_∞ says that there are infinitely many arches which pass above and below each vertex of \mathbf{A} . This guarantees that \mathbf{A} has no unbounded face. The rest of the construction consists of local operations only. So the resulting subgraph-rooted map $(\mathbf{M}, \mathbf{G}) = \Psi(w)$ is a locally finite subgraph-rooted map.

Also, by Proposition 5, we have $\mathbb{P}_\infty^{(p)}(\mathcal{W}_\infty) = 1$. It remains to see that (the extension) of Ψ is continuous on $\mathcal{W}' = \bigcup_n \mathcal{W}_n \cup \mathcal{W}_\infty$. Let $w^{(n)}, w \in \mathcal{W}'$ so that $w^{(n)} \rightarrow w$ for D_{loc} . If w is finite, there is nothing to prove. Otherwise, let $(\mathbf{M}, \mathbf{G}) = \Psi(w)$ and consider a ball B of finite radius r around the root in the map \mathbf{M} . By locality of the mapping $\Delta \rightsquigarrow \mathbf{M}$, the ball B can be determined by a ball B' of finite radius r' (which may depend on \mathbf{M}) in Δ . But each triangle in Δ corresponds to a letter in the word, so there exists r'' (which may depend on \mathbf{M}) such that if $w_{[-r'', r'']}^{(n)} = w_{[-r'', r'']}$ then the balls of radius r' in Δ coincide. This proves that Ψ is continuous on \mathcal{W}_∞ . \square




Figure 6: An example of an infinite arch graph, and the associated process $(X_k)_{k \in \mathbb{Z}}$. We shifted the arch graph horizontally by $1/2$ relative to the graph of X , since the time k for the process X naturally corresponds to the interval $[k - 1, k]$ in the arch graph.

4.2 Proof of Theorem 1

Combining Proposition 4, Theorem 8 and Proposition 9 yields the convergence statement of the theorem. The self-duality of the infinite cFK random map follows from the finite self-duality. It remains to show that the infinite cFK random maps are almost surely one-ended and recurrent for all q , and that their laws are mutually singular for different values of q .

One-endedness. Recall that a graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ is said to be one-ended if for any finite subset of vertices \mathbf{U} , $\mathbf{V} \setminus \mathbf{U}$ has exactly one infinite connected component. We will prove that for a word $w \in \mathcal{W}_\infty$, $\Psi(w)$ is one-ended. Let \mathbf{A} (resp. $\mathbf{\Delta}$) be the arch graph (resp. triangulation) associated to w , and let $(\mathbf{M}, \mathbf{G}) = \Psi(w)$. By the second condition in the definition of \mathcal{W}_∞ (see (20)), there exist arches that connect vertices on the left of w_{-R} to vertices on the right of w_R for any finite number R . Therefore the arch graph \mathbf{A} is one-ended. It is then an easy exercise to deduce from this that the triangulation $\mathbf{\Delta}$ and then the map \mathbf{M} are also one-ended.

Recurrence. To prove the recurrence of \mathbf{M} we use the general criterion established by Gurel-Gurevich and Nachmias [20]. Notice first that under $\mathbb{Q}_n^{(q)}$, the random maps are *uniformly rooted*, that is, conditionally on the map, the root vertex ρ is chosen with probability proportional to its degree. By [20] it thus suffices to check that the distribution of $\deg(\rho)$ has an exponential tail. For this we claim that the variable N_0 studied in Lemma 6 exactly corresponds to the degree of the root in an infinite cFK random map. From the construction of the hamburger-cheeseburger bijection, we see that the vertices of the map \mathbf{M} corresponds to the faces of the arch graph *in the upper half plane*. In particular, the root vertex ρ corresponds to the face *above* the interval $[-1, 0]$, and $\deg(\rho)$ is the number of unit intervals on the real axis which are also on the boundary of this face. On the other hand, X_k is the *net number* of arches that one *enters* to get from the face above $[-1, 0]$ to the face above $[k - 1, k]$, see Fig. 6. So N_0 exactly counts the above number of intervals.

Mutual singularity. The basic idea is to construct a measurable function f on infinite rooted planar maps such that $f(\mathbf{M})$ is almost surely constant for each $q \in [0, \infty]$, and that $f(\mathbf{M}) = f(q)$ is an injective function of q . One such function can be defined as follows. Consider the simple random walk on the map \mathbf{M} . We regard it as a sequence of oriented edges $(\vec{E}_k, k \geq 0)$ such that \vec{E}_0 is the the root edge (i.e. the edge on the left of the root corner), and such that \vec{E}_{k+1} starts at the vertex where \vec{E}_k ends. An oriented edge \vec{e} is *pending* if the starting point of \vec{e} has degree 1. Let

$$f(\mathbf{M}) = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{1}_{\{\vec{E}_k \text{ is pending}\}}$$

From the ergodicity result in the appendix and Birkhoff's ergodic theorem, it follows that $f(\mathbf{M}) = \mathbb{Q}(\vec{E}_0 \text{ is pending})$ almost surely. Recall that \vec{E}_0 is the root edge of \mathbf{M} . A moment of look at Fig. 5 shows that \vec{E}_0 is a pending edge in $(\mathbf{M}, \mathbf{G}) = \Phi(w)$ if and only if $w_{-1} = \mathbf{a}$ and $w_0 \in \{\mathbf{A}, \mathbf{F}\}$. Therefore,

$$\begin{aligned} \mathbb{Q}(\vec{E}_0 \text{ is pending}) &= \mathbb{P}(w_{-1} = \mathbf{a} \text{ and } w_0 \in \{\mathbf{A}, \mathbf{F}\}) \\ &= \theta(\mathbf{a})(\theta(\mathbf{A}) + \theta(\mathbf{F})) = \frac{1+p}{16} = \frac{1}{16} \left(1 + \frac{\sqrt{q}}{2 + \sqrt{q}}\right) \end{aligned}$$

We see that $f(q) = f(\mathbf{M})$ is injective in q .

A Ergodicity of cFK random maps

Here we use the framework set up by Benjamini and Curien in [3] and we follow the exposition in [14, Sec 3.1]. Consider the space $\vec{\mathcal{M}}$ of all locally finite rooted maps endowed with a path $\vec{e} = (\vec{e}_n)_{n \geq 0}$ starting from the root edge. Let \vec{d}_{loc} be the natural local distance on $\vec{\mathcal{M}}$

$$\vec{d}_{\text{loc}}((\mathbf{M}, \vec{e}), (\mathbf{M}', \vec{e}')) = \inf \left\{ 2^{-R} \mid B_R(\mathbf{M}) = B_R(\mathbf{M}') \text{ and } \vec{e}_k = \vec{e}'_k \text{ for } 0 \leq k \leq R \right\}$$

and consider the associated Borel σ -algebra. If \mathbf{M} is a rooted map, let $Q_{\mathbf{M}}$ be the law of the simple random walk starting from the root edge of \mathbf{M} . We denote by $\mathbb{Q} \otimes Q$ the probability measure

$$\mathbb{Q} \otimes Q(A) = \int Q_{\mathbf{M}}(A_{\mathbf{M}}) \mathbb{Q}(d\mathbf{M})$$

on $\vec{\mathcal{M}}$, where $A_{\mathbf{M}} = \{(\vec{e}_n)_{n \geq 0} \mid (\mathbf{M}, (\vec{e}_n)_{n \geq 0}) \in A\}$. To simplify notation, we will write $Q_{\mathbf{M}}(A)$ instead of $Q_{\mathbf{M}}(A_{\mathbf{M}})$ in the sequel.

If \vec{e} is an oriented edge of \mathbf{M} , we write $\mathbf{M}|_{\vec{e}}$ for the map obtained by re-rooting \mathbf{M} at \vec{e} . Observe that if $(\mathbf{M}, (\vec{E}_n)_{n \geq 0})$ is a random variable of law $\mathbb{Q} \otimes Q$, then

$$\mathbf{M} = \mathbf{M}|_{\vec{E}_1}$$

in distribution. (Remark that the left-hand-side is the same as $\mathbf{M}|_{\vec{E}_0}$.) This is a consequence of the fact that \mathbf{M} is the local limit of a sequence of uniformly rooted finite maps. It follows that the shift operator

$$\tau : (\mathbf{M}, (\vec{e}_n)_{n \geq 0}) \mapsto (\mathbf{M}|_{\vec{e}_1}, (\vec{e}_{n+1})_{n \geq 0})$$

preserves the measure $\mathbb{Q} \otimes Q$.

Proposition 10. *The shift operator τ is ergodic for $\mathbb{Q} \otimes Q$.*

Proof. Let A be a τ -invariant event for $\mathbb{Q} \otimes Q$, that is, $\mathbb{Q} \otimes Q(A \Delta \tau^{-1}(A)) = 0$, where Δ denotes the symmetric difference: $A \Delta B = (A \setminus B) \cup (B \setminus A)$. We do the proof in three steps:

(a) \mathbb{Q} -almost surely, $Q_{\mathbf{M}}(A) \in \{0, 1\}$.

Recall that the simple random walk on \mathbf{M} is \mathbb{Q} -almost surely recurrent. Therefore, it can be decomposed into an i.i.d. sequence of excursions from the root edge. An event invariant by τ is also invariant by the shift operator associated to this i.i.d. sequence. Thus (a) follows from Kolmogorov's zero-one law.

(b) \mathbb{Q} -almost surely, the value of $Q_{\mathbf{M}}(A)$ is invariant under any re-rooting of the map \mathbf{M} .

First let us fix a rooted map \mathbf{M} . Let x be the vertex to which the root edge points. Denote by d the degree of x , and $\vec{\rho}_1, \dots, \vec{\rho}_d$ the d oriented edges that point away from x . We deduce from the Markov property of the simple random walk that

$$\begin{aligned} Q_{\mathbf{M}}(\tau^{-1}A) &= Q_{\mathbf{M}}\{(\vec{e}_n)_{n \geq 0} \mid (\mathbf{M}|_{\vec{e}_1}, (\vec{e}_{n+1})_{n \geq 0}) \in A\} \\ &= \sum_{k=1}^d Q_{\mathbf{M}}\{(\vec{e}_n)_{n \geq 0} \mid (\mathbf{M}|_{\vec{e}_1}, (\vec{e}_{n+1})_{n \geq 0}) \in A \text{ and } \vec{e}_1 = \vec{\rho}_k\} \\ &= \sum_{k=1}^d \frac{1}{d} Q_{\mathbf{M}|_{\vec{\rho}_k}}\{(\vec{e}_n)_{n \geq 0} \mid (\mathbf{M}|_{\vec{\rho}_k}, (\vec{e}_n)_{n \geq 0}) \in A\} = \frac{1}{d} \sum_{k=1}^d Q_{\mathbf{M}|_{\vec{\rho}_k}}(A) \end{aligned}$$

From

$$\int |Q_{\mathbf{M}}(A) - Q_{\mathbf{M}}(\tau^{-1}A)| \mathbb{Q}(d\mathbf{M}) \leq \mathbb{Q} \otimes Q(A \Delta \tau^{-1}A) = 0$$

we deduce that \mathbb{Q} -almost surely,

$$Q_{\mathbf{M}}(A) = \frac{1}{d} \sum_{k=1}^d Q_{\mathbf{M}|_{\vec{\rho}_k}}(A)$$

But according to (a), $Q_{\mathbf{M}}(A)$ and $Q_{\mathbf{M}|_{\vec{\rho}_k}}(A)$ ($k = 1, \dots, d$) are either 0 or 1. Thus \mathbb{Q} -almost surely, $Q_{\mathbf{M}}(A) = 1$ if and only if for all $k = 1, \dots, d$, $Q_{\mathbf{M}|_{\vec{\rho}_k}}(A) = 1$. In other words, the value of $Q_{\mathbf{M}}(A)$ is unchanged when re-rooting at a neighbor of the root edge. But since the maps are connected, we obtain (b) by iterating the above argument.

- (c) If an event \hat{A} on the space of (locally finite) rooted maps is \mathbb{Q} -almost surely invariant under re-rooting, then $\mathbb{Q}(\hat{A}) \in \{0, 1\}$.

Consider the measure-preserving mapping $\hat{\Phi} : w \mapsto \mathbf{M}$, where $(\mathbf{M}, \mathbf{G}) = \Phi(w)$ for some subgraph \mathbf{G} . Via this mapping, a translation of indices in a word w give rise to a re-rooting of the corresponding map \mathbf{M} . Therefore if an event \hat{A} is \mathbb{Q} -almost surely invariant under re-rooting, then $\hat{\Phi}^{-1}\hat{A}$ is \mathbb{P} -almost surely invariant under translation of the indices. But, under \mathbb{P} the letters of w are i.i.d. random variables, so we have $\mathbb{Q}(\hat{A}) = \mathbb{P}(\hat{\Phi}^{-1}\hat{A}) \in \{0, 1\}$.

Finally, considering $\hat{A} = \{\mathbf{M} : Q_{\mathbf{M}}(A) = 1\}$ shows that $\mathbb{Q} \otimes Q(A) \in \{0, 1\}$, as desired. \square

Remark. Only the step (c) of the proof uses specific features of the cFK random maps, namely, their representation by an i.i.d. sequence of letters. For any infinite random map whose law is stationary under τ , the proof of (a) and (b) goes through provided that the random map is almost surely recurrent. (Note that the proof of (b) depends on (a).) The case of almost surely transient random maps was treated in [14, Prop. 10]. There, (a) was proved under the following reversibility condition: if \overleftarrow{e} is the root edge of \mathbf{M} , oriented in the opposite direction, then

$$\mathbf{M} = \mathbf{M}|_{\overleftarrow{e}}$$

in distribution. And (b) was replaced by the following variant, which follows directly from the transience of the map.

- (b') Almost surely, $Q_{\mathbf{M}}(A)$ is unchanged by any finite modification of the map \mathbf{M} .

References

- [1] J. Ambjørn, B. Durhuus, and T. Jonsson. *Quantum geometry: a statistical field theory approach*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, Cambridge, 1997.
- [2] O. Angel and O. Schramm. Uniform infinite planar triangulations. *Comm. Math. Phys.*, 241(2-3):191–213, 2003. arXiv:math/0207153.
- [3] I. Benjamini and N. Curien. Ergodic theory on stationary random graphs. *Electron. J. Probab.*, 17:no. 93, 20 pp. (electronic), 2012. arXiv:1011.2526
- [4] I. Benjamini and O. Schramm. Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, 6:no. 23, 13 pp. (electronic), 2001.
- [5] N. Berestycki, B. Laslier, and G. Ray. Critical exponents on fortuin–kastelyn weighted planar maps. *Preprint*, 2015. arXiv:1502.00450.
- [6] O. Bernardi. Bijective counting of tree-rooted maps and shuffles of parenthesis systems. *Electron. J. Combin.*, 14(1):Research Paper 9, 36 pp. (electronic), 2007.
- [7] O. Bernardi and M. Bousquet-Mélou. Counting colored planar maps: algebraicity results. *J. Combin. Theory Ser. B*, 101(5):315–377, 2011.
- [8] J. E. Björnberg and S. Ö. Stefánsson. Recurrence of bipartite planar maps. *Electron. J. Probab.*, 19:no. 31, 40, 2014. arXiv:1311.0178.
- [9] G. Borot, J. Bouttier, and E. Guitter. Loop models on random maps via nested loops: the case of domain symmetry breaking and application to the Potts model. *J. Phys. A*, 45(49):494017, 35, 2012. arXiv:1207.4878.
- [10] G. Borot, J. Bouttier, and E. Guitter. More on the $O(n)$ model on random maps via nested loops: loops with bending energy. *J. Phys. A*, 45(27):275206, 32, 2012. arXiv:1202.5521.
- [11] G. Borot, J. Bouttier, and E. Guitter. A recursive approach to the $O(n)$ model on random maps via nested loops. *J. Phys. A*, 45(4):045002, 38, 2012. arXiv:1106.0153.
- [12] J. Bouttier, P. Di Francesco, and E. Guitter. Planar maps as labeled mobiles. *Electron. J. Combin.*, 11(1):Research Paper 69, 27, 2004. arXiv:math/0405099.
- [13] P. Chassaing and B. Durhuus. Local limit of labeled trees and expected volume growth in a random quadrangulation. *Ann. Probab.*, 34(3):879–917, 2006. arXiv:math/0311532.
- [14] N. Curien. Planar stochastic hyperbolic triangulations. *Probab. Theory Related Fields*, 2015. arXiv:1401.3297.
- [15] N. Curien, L. Ménard, and G. Miermont. A view from infinity of the uniform infinite planar quadrangulation. *ALEA Lat. Am. J. Probab. Math. Stat.*, 10(1):45–88, 2013. arXiv:1201.1052.
- [16] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume 38 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, second edition, 1998.

- [17] B. Eynard and C. Kristjansen. Exact solution of the $O(n)$ model on a random lattice. *Nuclear Phys. B*, 455(3):577–618, 1995.
- [18] G. Grimmett. *The random-cluster model*, volume 333 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2006.
- [19] A. Guionnet, V. F. R. Jones, D. Shlyakhtenko, and P. Zinn-Justin. Loop models, random matrices and planar algebras. *Comm. Math. Phys.*, 316(1):45–97, 2012.
- [20] O. Gurel-Gurevich and A. Nachmias. Recurrence of planar graph limits. *Ann. of Math. (2)*, 177(2):761–781, 2013. arXiv:1206.0707.
- [21] E. Gwynne, C. Mao, and X. Sun. Scaling limits for the critical fortuin-kasteleyn model on a random planar map I: cone times. *Preprint*, 2015. arXiv:1502.00546.
- [22] E. Gwynne, and X. Sun. Scaling limits for the critical fortuin-kasteleyn model on a random planar map II: local estimates and empty reduced word exponent. *Preprint*, 2015. arXiv:1505.03375.
- [23] E. Gwynne, and X. Sun. Scaling limits for the critical fortuin-kasteleyn model on a random planar map III: finite volume case. *Preprint*, 2015. arXiv:1510.06346.
- [24] I. K. Kostov. $O(n)$ vector model on a planar random lattice: spectrum of anomalous dimensions. *Modern Phys. Lett. A*, 4(3):217–226, 1989.
- [25] M. Krikun. Local structure of random quadrangulations. *Preprint*, 2006. arXiv:math/0512304.
- [26] L. Ménard. The two uniform infinite quadrangulations of the plane have the same law. *Ann. Inst. Henri Poincaré Probab. Stat.*, 46(1):190–208, 2010. arXiv:0812.0965.
- [27] G. Miermont. Aspects of random maps. Lecture notes of the 2014 Saint-Flour Probability Summer School, <http://perso.ens-lyon.fr/gregory.miermont/coursSaint-Flour.pdf>.
- [28] G. Schaeffer. *Conjugaison d’arbres et cartes combinatoires aléatoires*. PhD thesis, Université de Bordeaux 1, 1998.
- [29] S. Sheffield. Quantum gravity and inventory accumulation. *Preprint*, 2011. arXiv:1108.2241.

Linxiao Chen, Institut de Physique Théorique, Université Paris Saclay, CEA, CNRS, F-91191 Gif-sur-Yvette, France.
 Laboratoire de Mathématiques d’Orsay, Univ. Paris-Sud, CNRS, Université Paris-Saclay, 91405 Orsay, France.